

Adaptive multiresolution analysis based on anisotropic triangulations

Albert Cohen, Nira Dyn, Frédéric Hecht
and Jean-Marie Mirebeau ^{*}

January 10, 2011

Abstract

A simple greedy refinement procedure for the generation of data-adapted triangulations is proposed and studied. Given a function f of two variables, the algorithm produces a hierarchy of triangulations $(\mathcal{D}_j)_{j \geq 0}$ and piecewise polynomial approximations of f on these triangulations. The refinement procedure consists in bisecting a triangle T in a direction which is chosen so as to minimize the local approximation error in some prescribed norm between f and its piecewise polynomial approximation after T is bisected. The hierarchical structure allows us to derive various approximation tools such as multiresolution analysis, wavelet bases, adaptive triangulations based either on greedy or optimal CART trees, as well as a simple encoding of the corresponding triangulations. We give a general proof of convergence in the L^p norm of all these approximations. Numerical tests performed in the case of piecewise linear approximation of functions with analytic expressions or of numerical images illustrate the fact that the refinement procedure generates triangles with an optimal aspect ratio (which is dictated by the local Hessian of f in case of C^2 functions).

1 Introduction

Approximation by piecewise polynomial functions is a standard procedure which occurs in various applications. In some of them such as terrain data simplification or image compression, the function to be approximated might be fully known, while it might be only partially known or fully unknown in other applications such as denoising, statistical learning or in the finite element discretization of PDE's.

In all these applications, one usually makes the distinction between *uniform* and *adaptive* approximation. In the uniform case, the domain of interest is decomposed into a partition where all elements have comparable shape and size, while these attributes are allowed to vary strongly in the adaptive case. In the context of adaptive triangulations, another important distinction is between *isotropic* and *anisotropic* triangulations. In the first case the triangles satisfy a condition which guarantees that they do not differ too much from equilateral triangles. This can either be stated in terms of a minimal value $\theta_0 > 0$ for every angle, or by a uniform bound on the aspect ratio

$$\rho_T := \frac{h_T}{r_T}$$

of each triangle T where h_T and r_T respectively denote the diameter of T and of its largest inscribed disc. In the second case, which is in the scope of the present paper, the aspect ratio is allowed to be arbitrarily large, i.e. long and thin triangles are allowed. In summary, adaptive and anisotropic triangulations mean that we do not fix any constraint on the size and shape of the triangles.

Given a function f and a norm $\|\cdot\|_X$ of interest, we can formulate the problem of finding the *optimal triangulation* for f in the X -norm in two related forms:

^{*}This research was supported by the P2R French-Israeli program “Nonlinear multiscale methods - applications to image and terrain data”

- For a given N find a triangulation \mathcal{T}_N with N triangles and a piecewise polynomial function f_N (of some fixed degree) on \mathcal{T}_N such that $\|f - f_N\|_X$ is minimized.
- For a given $\varepsilon > 0$ find a triangulation \mathcal{T}_N with minimal number of triangles N and a piecewise polynomial function f_N such that $\|f - f_N\|_X \leq \varepsilon$.

In this paper X will be the L^p norm for some arbitrary $1 \leq p \leq \infty$. The exact solution to such problems is usually out of reach both analytically and algorithmically: even when restricting the search of the vertices to a finite grid, the number of possible triangulations has combinatorial complexity and an exhaustive search is therefore prohibitive.

Concrete mesh generation algorithms have been developed in order to generate in reasonable time triangulations which are “close” to the above described optimal trade-off between error and complexity. They are typically governed by two intuitively desirable features:

1. The triangulation should *equidistribute* the local approximation error between each triangle. This rationale is typically used in local mesh refinement algorithms for numerical PDE’s [23]: a triangle is refined when the local approximation error (estimated by an a-posteriori error indicator) is large.
2. In the case of anisotropic meshes, the local aspect ratio should in addition be optimally adapted to the approximated function f . In the case of piecewise linear approximation, this is achieved by imposing that the triangles are isotropic with respect to a distorted metric induced by the Hessian $d^2 f$. We refer in particular to [7] where this task is executed using Delaunay mesh generation techniques.

While these last algorithms fastly produce anisotropic meshes which are naturally adapted to the approximated function, they suffer from two intrinsic limitations:

1. They are based on the evaluation of the Hessian $d^2 f$, and therefore do not in principle apply to arbitrary functions $f \in L^p(\Omega)$ for $1 \leq p \leq \infty$ or to noisy data.
2. They are non-hierarchical: for $N > M$, the triangulation \mathcal{T}_N is not a refinement of \mathcal{T}_M .

One way to circumvent the first limitation is to regularize the function f , either by projection onto a finite element space or by convolution by a mollifier. However this raises the additional problem of appropriately tuning the amount of smoothing, in particular depending on the noise level in f .

The need for hierarchical triangulations is critical in the construction of wavelet bases, which play an important role in applications to image and terrain data processing, in particular data compression [11]. In such applications, the multilevel structure is also of key use for the fast encoding of the information. Hierarchy is also useful in the design of optimally converging adaptive methods for PDE’s [17, 21, 5, 22]. However, all these developments are so far mostly restricted to isotropic refinement methods. Let us mention that hierarchical and anisotropic triangulations have been investigated in [18], yet in this work the triangulations are *fixed in advance* and therefore generally not adapted to the approximated function.

A natural objective is therefore to design adaptive algorithmic techniques that combine hierarchy and anisotropy, and that apply to any function $f \in L^p(\Omega)$, without any need for regularization.

In this paper we propose and study a simple *greedy refinement procedure* that achieves this goal: starting from an initial triangulation \mathcal{D}_0 , the procedure bisects every triangle from one of its vertices to the mid-point of the opposite segment. The choice of the vertex is typically the one which minimizes the new approximation error after bisection among the three options.

Surprisingly, it turns out that - in the case of piecewise linear approximation - this elementary strategy tends to generate anisotropic triangles with optimal aspect ratio. This fact is rigorously proved in [12] which establishes optimal error estimates for the approximation of *smooth and convex functions* $f \in C^2$, by adaptive triangulations \mathcal{T}_N with N triangles. These triangulations are obtained by consecutively applying the refinement procedure to the triangle of maximal error. The estimates in [12] are of the form

$$\|f - f_N\|_{L^p} \leq CN^{-1} \|\sqrt{|\det(d^2 f)|}\|_{L^\tau}, \quad \frac{1}{\tau} = \frac{1}{p} + 1. \quad (1.1)$$

and were already established in [10, 4] for functions which are not necessarily convex, however based on triangulations which are non-hierarchical and based on the evaluation of d^2f . Note that (1.1) improves on the estimate

$$\|f - f_N\|_{L^p} \leq CN^{-1}\|d^2f\|_{L^\tau}, \quad \frac{1}{\tau} = \frac{1}{p} + 1. \quad (1.2)$$

which can be established (see §2 in [12]) for adaptive triangulations with isotropic triangles, and which itselfs improves on the classical estimate

$$\|f - f_N\|_{L^p} \leq CN^{-1}\|d^2f\|_{L^p}, \quad (1.3)$$

which is known to hold for uniform triangulations.

The main objective of the present paper is to introduce the refinement procedure as well as several approximation methods based on it, and to study their convergence for *an arbitrary function* $f \in L^p$. In §2, we introduce notation that serves to describe the refinement procedure and define the anisotropic hierarchy of triangulations $(\mathcal{D}_j)_{j \geq 0}$. We show how this general framework can be used to derive adaptive approximations of f either by triangulations based on greedy or optimal trees, or by wavelet thresholding. In §3, we show that as defined, the approximations produced by the refinement procedure may fail to converge for certain $f \in L^p$ and show how to modify the procedure so that convergence holds for any arbitrary $f \in L^p$. We finally present in §4 some numerical tests which illustrate the optimal mesh adaptation, in the case of piecewise linear elements, when the refinement procedure is applied either to synthetic functions or to numerical images.

2 An adaptive and anisotropic multiresolution framework

2.1 The refinement procedure

Our refinement procedure is based on a local approximation operator \mathcal{A}_T acting from $L^p(T)$ onto Π_m - the space of polynomials of total degree less or equal to m . Here, the parameters $m \geq 0$ and $1 \leq p \leq \infty$ are arbitrary but fixed. For a generic triangle T , we denote by (a, b, c) its edge vectors oriented in clockwise or anticlockwise direction so that

$$a + b + c = 0.$$

We define the local L^p approximation error

$$e_T(f)_p := \|f - \mathcal{A}_T f\|_{L^p(T)}.$$

The most natural choice for \mathcal{A}_T is the operator \mathcal{B}_T of best $L^p(T)$ approximation which is defined by

$$\|f - \mathcal{B}_T f\|_{L^p(T)} = \min_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}.$$

However this operator is non-linear and not easy to compute when $p \neq 2$. In practice, one prefers to use an operator which is easier to compute, yet nearly optimal in the sense that

$$\|f - \mathcal{A}_T f\|_{L^p(T)} \leq C \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}. \quad (2.4)$$

with C a Lebesgue constant independent of f and T . Two particularly simple admissible choices of approximation operators are the following:

1. $\mathcal{A}_T = P_T$, the $L^2(T)$ -orthogonal projection onto Π_m , defined by $P_T f \in \Pi_m$ such that $\int_T (f - P_T f) \pi = 0$ for all $\pi \in \Pi_m$. This operator has finite Lebesgue constant for all p , with $C = 1$ when $p = 2$ and $C > 1$ otherwise.
2. $\mathcal{A}_T = I_T$, the local interpolation operator which is defined by $I_T f \in \Pi_m$ such that $I_T f(\gamma) = f(\gamma)$ for all $\gamma \in \Sigma := \{\sum \frac{k_i}{m} v_i ; k_i \in \mathbb{N}, \sum k_i = m\}$ where $\{v_1, v_2, v_3\}$ are the vertices of T (in the case $m = 0$ we can take for Σ the barycenter of T). This operator is only defined on continuous functions and has Lebesgue constant $C > 1$ in the L^∞ norm.

All our results are simultaneously valid when \mathcal{A}_T is either P_T or I_T (in the case where $p = \infty$), or any linear operator that fulfills the continuity assumption (2.4).

Given a target function, our refinement procedure defines by induction a hierarchy of nested triangulations $(\mathcal{D}_j)_{j \geq 0}$ with $\#(\mathcal{D}_j) = 2^j \#(\mathcal{D}_0)$. The procedure starts from the coarse triangulation \mathcal{D}_0 of Ω , which is fixed independently of f . When $\Omega = [0, 1]^2$ we may split it into two symmetric triangles so that $\#(\mathcal{D}_0) = 2$. For every $T \in \mathcal{D}_j$, we split T into two sub-triangles of equal area by bisection from one of its three vertices towards the mid-point of the opposite edge $e \in \{a, b, c\}$. We denote by T_e^1 and T_e^2 the two resulting triangles. The choice of $e \in \{a, b, c\}$ is made according to a *refinement rule* that selects this edge depending on the properties of f . We denote by \mathcal{R} this refinement rule, which can therefore be viewed as a mapping

$$\mathcal{R} : (f, T) \mapsto e.$$

We thus obtain two *children* of T corresponding to the choice e . \mathcal{D}_{j+1} is the triangulation consisting of all such pairs corresponding to all $T \in \mathcal{D}_j$.

In this paper, we consider refinement rules where the selected edge e minimizes a *decision function* $e \mapsto d_T(e, f)$ among $\{a, b, c\}$. We refer to such rules as *greedy refinement rules*. A more elaborate type of refinement rule is also considered in §3.3.

The role of the decision function is to drive the generation of anisotropic triangles according to the local properties of f , in contrast to simpler procedures such as *newest vertex bisection* (i.e. split T from the most recently created vertex) which is independent of f and generates triangulations with isotropic shape constraint.

Therefore, the choice of $d_T(e, f)$ is critical in order to obtain triangles with an optimal aspect ratio. The most natural choice corresponds to the optimal split

$$d_T(e, f) = e_{T_e^1}(f)_p^p + e_{T_e^2}(f)_p^p,$$

i.e. choose the edge that minimizes the resulting L^p error after bisection. It is proved in [12] in the case of piecewise linear approximation, that when f is a C^2 function which is strictly convex or concave the refinement rule based on the decision function

$$d_T(e, f) = \|f - I_{T_e^1} f\|_{L^1(T_e^1)} + \|f - I_{T_e^2} f\|_{L^1(T_e^2)}. \quad (2.5)$$

generates triangles which tend to have an optimal aspect ratio, locally adapted to the Hessian $d^2 f$. This aspect ratio is independent of the L^p norm in which one wants to minimize the error between f and its piecewise affine approximation.

Remark 2.1 *If the minimizer e is not unique, we may choose it among the multiple minimizers either randomly or according to some prescribed ordering of the edges (for example the largest coordinate pair of the opposite vertex in lexicographical order).*

Remark 2.2 *The triangulations \mathcal{D}_j which are generated by the greedy procedure are in general non-conforming, i.e. exhibit hanging nodes. This is not problematic in the present setting since we consider approximation in the L^p norm which does not require global continuity of the piecewise polynomial functions.*

The refinement rule \mathcal{R} defines a multiresolution framework. For a given $f \in L^p(\Omega)$ and any triangle T we denote by

$$\mathcal{C}(T) := \{T_1, T_2\},$$

the *children* of T which are the two triangles obtained by splitting T based on the prescribed decision function $d_T(e, f)$. We also say that T is the parent of T_1 and T_2 and write

$$T = \mathcal{P}(T_1) = \mathcal{P}(T_2).$$

Note that

$$\mathcal{D}_j := \cup_{T \in \mathcal{D}_{j-1}} \mathcal{C}(T).$$

We also define

$$\mathcal{D} := \cup_{j \geq 0} \mathcal{D}_j,$$

which has the structure of *an infinite binary tree*. Note that \mathcal{D}_j depends on f (except for $j = 0$) and on the refinement rule \mathcal{R} , and thus \mathcal{D} also depends on f and \mathcal{R} :

$$\mathcal{D}_j = \mathcal{D}_j(f, \mathcal{R}) \text{ and } \mathcal{D} = \mathcal{D}(f, \mathcal{R}).$$

For notational simplicity, we sometimes omit the dependence in f and \mathcal{R} when there is no possible ambiguity.

2.2 Adaptive tree-based triangulations

A first application of the multiresolution framework is the design of adaptive anisotropic triangulations \mathcal{T}_N for piecewise polynomial approximation, by a *greedy tree algorithm*. For any finite sub-tree $\mathcal{S} \subset \mathcal{D}$, we denote by

$$\mathcal{L}(\mathcal{S}) := \{T \in \mathcal{S} \text{ s.t. } \mathcal{C}(T) \notin \mathcal{S}\}$$

its *leaves* which form a partition of Ω . We also denote by

$$\mathcal{I}(\mathcal{S}) := \mathcal{S} \setminus \mathcal{L}(\mathcal{S}),$$

its *inner nodes*. Note that *any* finite partition of Ω by elements of \mathcal{D} is the set of leaves of a finite sub-tree. One easily checks that

$$\#(\mathcal{S}) = 2\#(\mathcal{L}(\mathcal{S})) - N_0.$$

For each N , the greedy tree algorithm defines a finite sub-tree \mathcal{S}_N of \mathcal{D} which grows from $\mathcal{S}_{N_0} := \mathcal{D}_0 = \mathcal{T}_{N_0}$, by adding to \mathcal{S}_{N-1} the two children of the triangle T_{N-1}^* which maximizes the local L^p -error $e_T(f)_p$ over all triangles in \mathcal{T}_{N-1} .

The adaptive partition \mathcal{T}_N associated with the greedy algorithm is defined by

$$\mathcal{T}_N := \mathcal{L}(\mathcal{S}_N).$$

Similarly to \mathcal{D} , the triangulation \mathcal{T}_N depends on f and on the refinement rule \mathcal{R} , but also on p and on the choice of the approximation operator \mathcal{A}_T . We denote by f_N the piecewise polynomial approximation to f which is defined as $\mathcal{A}_T f$ on each $T \in \mathcal{T}_N$. The global L^p approximation error is thus given by

$$\|f - f_N\|_{L^p} = \|(e_T(f)_p)\|_{\ell^p(\mathcal{T}_N)}.$$

Stopping criterions for the algorithm can be defined in various ways:

- Number of triangles: stop once a prescribed N is attained.
- Local error: stop once $e_T(f)_p \leq \varepsilon$ for all $T \in \mathcal{T}_N$, for some prescribed $\varepsilon > 0$.
- Global error: stop once $\|f - f_N\|_{L^p} \leq \varepsilon$ for some prescribed $\varepsilon > 0$.

Remark 2.3 *The role of the triangle selection based on the largest $e_T(f)_p$ is to equidistribute the local L^p error, a feature which is desirable when we want to approximate f in $L^p(\Omega)$ with the smallest number of triangles. However, it should be well understood that the refinement rule may still be chosen based on a decision function defined by approximation errors in norms that differ from L^p . In particular, as explained earlier, the decision function (2.5) generates triangles which tend to have an optimal aspect ratio, locally adapted to the Hessian $d^2 f$ when f is strictly convex or concave, and this aspect ratio is independent of the L^p norm in which one wants to minimize the error between f and its piecewise affine approximation.*

The greedy algorithm is one particular way of deriving an adaptive triangulation for f within the multiresolution framework defined by the infinite tree \mathcal{D} . An interesting alternative is to build adaptive triangulations within \mathcal{D} which offer an *optimal trade-off between error and complexity*. This can be done when $1 \leq p < \infty$ by solving the minimization problem

$$\min_{\mathcal{S}} \left\{ \sum_{T \in \mathcal{L}(\mathcal{S})} e_T(f)_p^p + \lambda \#(\mathcal{S}) \right\} \quad (2.6)$$

among all finite trees, for some fixed $\lambda > 0$. In this approach, we do not directly control the number of triangles which depends on the penalty parameter λ . However, it is immediate to see that if $N = N(\lambda)$ is the cardinality of $\mathcal{T}_N^* = \mathcal{L}(\mathcal{S}^*)$ where \mathcal{S}^* is the minimizing tree, then \mathcal{T}_N^* minimizes the L^p approximation error

$$\mathcal{T}_N^* := \operatorname{Argmin}_{\#(\mathcal{T}) \leq N} \sum_{T \in \mathcal{T}} e_T(f)_p^p,$$

where the minimum is taken among all partitions \mathcal{T} of Ω within \mathcal{D} of cardinality less than or equal to N .

Due to the additive structure of the error term, the minimization problem (2.6) can be performed in fast computational time using an optimal pruning algorithm of CART type, see [8, 16]. In the case $p = \infty$ the associated minimization problem

$$\min_{\mathcal{S}} \left\{ \sup_{T \in \mathcal{L}(\mathcal{S})} e_T(f)_\infty + \lambda \#(\mathcal{S}) \right\}, \quad (2.7)$$

can also be solved by a similar fast algorithm. It is obvious that this method improves over the greedy tree algorithm: if N is the cardinality of the triangulation resulting from the minimization in (2.6) and f_N^* the corresponding piecewise polynomial approximation of f associated with this triangulation, we have

$$\|f - f_N^*\|_{L^p} \leq \|f - f_N\|_{L^p},$$

where f_N is built by the greedy tree algorithm.

2.3 Anisotropic wavelets

The multiresolution framework allows us to introduce the piecewise polynomial multiresolution spaces

$$V_j = V_j(f, \mathcal{R}) := \{g \text{ s.t. } g|_T \in \Pi_m, T \in \mathcal{D}_j\},$$

which depend on f and on the refinement rule \mathcal{R} . These spaces are nested and we denote by

$$V = V(f, \mathcal{R}) = \cup_{j \geq 0} V_j(f, \mathcal{R}),$$

their union. For notational simplicity, we sometimes omit the dependence in f and \mathcal{R} when there is no possible ambiguity.

The V_j spaces may be used to construct wavelet bases, following the approach introduced in [2] and that we describe in our present setting.

The space V_j is equipped with an orthonormal *scaling function basis*:

$$\varphi_T^i, \quad i = 1, \dots, \frac{1}{2}(m+1)(m+2), \quad T \in \mathcal{D}_j,$$

where the φ_T^i for $i = 1, \dots, \frac{1}{2}(m+1)(m+2)$ are supported in T and constitute an orthonormal basis of Π_m in the sense of $L^2(T)$ for each $T \in \mathcal{D}$. There are several possible choices for such a basis. In the particular case where $m = 1$, a simple one is to take for T with vertices (v_1, v_2, v_3) ,

$$\varphi_T^i(v_i) = |T|^{-1/2}\sqrt{3} \text{ and } \varphi_T^i(v_j) = -|T|^{-1/2}\sqrt{3}, \quad j \neq i.$$

We denote by P_j the orthogonal projection onto V_j :

$$P_j g := \sum_{T \in \mathcal{D}_j} \sum_i \langle g, \varphi_T^i \rangle \varphi_T^i.$$

We next introduce for each $T \in \mathcal{D}_j$ a set of *wavelets*

$$\psi_T^i, \quad i = 1, \dots, \frac{1}{2}(m+1)(m+2),$$

which constitutes an orthonormal basis of the orthogonal complement of $\Pi_m(T)$ into $\Pi_m(T') \oplus \Pi_m(T'')$ with $\{T', T''\}$ the children of T . In the particular case where $m = 1$, a simple choice for such a basis

is as follows: if (v_1, v_2, v_3) and (w_1, w_2, w_3) denote the vertices of T' and T'' , with the convention that $v_1 = w_1$ and $v_2 = w_2$ denote the common vertices, the second one being the midpoint of the segment (v_3, w_3) (i.e. T has vertices (v_3, w_3, v_1)), then

$$\begin{aligned}\psi_T^1 &:= \frac{\varphi_{T'}^3 - \varphi_{T''}^3}{\sqrt{2}}, \\ \psi_T^2 &:= \frac{\varphi_{T'}^1 - \varphi_{T'}^2 - \varphi_{T''}^1 + \varphi_{T''}^2}{2}, \\ \psi_T^3 &:= \frac{\varphi_{T'}^1 - \varphi_{T'}^3 + \varphi_{T''}^1 - \varphi_{T''}^3}{2}.\end{aligned}$$

where $\varphi_{T'}^i$ and $\varphi_{T''}^i$ are the above defined scaling functions.

The family

$$\psi_T^i, \quad i = 1, \dots, \frac{1}{2}(m+1)(m+2), \quad T \in \mathcal{D}_j$$

constitutes an orthonormal basis of W_j , the L^2 -orthogonal complement of V^j in V^{j+1} . A multiscale orthonormal basis of V_j is given by

$$\{\varphi_T\}_{T \in \mathcal{D}_0} \cup \{\psi_T^i\}_{\substack{i=1, \dots, \frac{1}{2}(m+1)(m+2), \\ T \in \mathcal{D}_j, j=0, \dots, J-1}}$$

Letting J go to $+\infty$ we thus obtain that

$$\{\varphi_T\}_{T \in \mathcal{D}_0} \cup \{\psi_T^i\}_{\substack{i=1, \dots, \frac{1}{2}(m+1)(m+2), \\ T \in \mathcal{D}_j, j \geq 0}}$$

is an orthonormal basis of the space

$$V(f, \mathcal{R})_2 := \overline{V(f, \mathcal{R})}^{L^2(\Omega)} = \overline{\bigcup_{j \geq 0} V_j(f, \mathcal{R})}^{L^2(\Omega)}.$$

For the sake of notational simplicity, we rewrite this basis as

$$(\psi_\lambda)_{\lambda \in \Lambda},$$

Note that $V(f, \mathcal{R})$ is not necessarily dense in $L^2(\Omega)$ and so $V(f, \mathcal{R})_2$ is not always equal to $L^2(\Omega)$. Therefore, the expansion of an arbitrary function $g \in L^2(\Omega)$ in the above wavelet basis does not always converge towards g in $L^2(\Omega)$. The same remark holds for the L^p convergence of the wavelet expansion of an arbitrary function $g \in L^p(\Omega)$ (or $\mathcal{C}(\Omega)$ in the case $p = \infty$): L^p -convergence holds when the space

$$V(f, \mathcal{R})_p := \overline{V(f, \mathcal{R})}^{L^p(\Omega)}$$

coincides with $L^p(\Omega)$ (or contains $\mathcal{C}(\Omega)$ in the case $p = \infty$), since we have

$$\left\| f - \sum_{|\lambda| < j} d_\lambda \psi_\lambda \right\|_{L^p} = \|f - P_j f\|_{L^p} \leq C \inf_{g \in V_j} \|f - g\|_{L^p},$$

with C the Lebesgue constant in (2.4) for the orthogonal projector.

A sufficient condition for such a property to hold is obviously that the size of all triangles goes to 0 as the level j increases, i.e.

$$\lim_{j \rightarrow +\infty} \sup_{T \in \mathcal{D}_j} \text{diam}(T) = 0.$$

However, this condition might not hold for the hierarchy $(\mathcal{D}_j)_{j \geq 0}$ produced by the refinement procedure. On the other hand, the multiresolution approximation being intrinsically adapted to f , a more reasonable requirement is that the expansion of f converges towards f in $L^p(\Omega)$ when $f \in L^p(\Omega)$ (or $\mathcal{C}(\Omega)$ in the case $p = \infty$). This is equivalent to the property

$$f \in V(f, \mathcal{R})_p.$$

We may then define an adaptive approximation of f by thresholding its coefficients at some level $\varepsilon > 0$:

$$f_\varepsilon := \sum_{|f_\lambda| \geq \varepsilon} f_\lambda \psi_\lambda,$$

where $f_\lambda := \langle f, \psi_\lambda \rangle$. When measuring the error in the L^p norm, a more natural choice is to perform thresholding on the component of the expansion measured in this norm, defining therefore

$$f_\varepsilon := \sum_{\|f_\lambda \psi_\lambda\|_{L^p} \geq \varepsilon} f_\lambda \psi_\lambda.$$

We shall next see that the condition $f \in V(f, \mathcal{R})_p$ also ensures the convergence of the tree-based adaptive approximations f_N and f_N^* towards f in $L^p(\Omega)$. We shall also see that this condition may not hold for certain functions f , but that this difficulty can be circumvented by a modification of the refinement procedure.

3 Convergence analysis

3.1 A convergence criterion

The following result relates the convergence towards f of its approximations by projection onto the spaces $V_j(f, \mathcal{R})$, greedy and optimal tree algorithms, and wavelet thresholding. This result is valid for *any* refinement rule \mathcal{R} .

Theorem 3.1 *Let $\mathcal{R} : (f, T) \mapsto e$ be an arbitrary refinement rule and let $f \in L^p(\Omega)$. The following statements are equivalent:*

- (i) $f \in V(f, \mathcal{R})_p$.
- (ii) *The greedy tree approximation converges: $\lim_{N \rightarrow +\infty} \|f - f_N\|_{L^p} = 0$.*
- (iii) *The optimal tree approximation converges: $\lim_{N \rightarrow +\infty} \|f - f_N^*\|_{L^p} = 0$.*

In the case $p = 2$, they are also equivalent to:

- (iv) *The thresholding approximation converges: $\lim_{\varepsilon \rightarrow 0} \|f - f_\varepsilon\|_{L^2} = 0$.*

Proof: Clearly, (ii) implies (iii) since $\|f - f_N^*\|_{L^p} \leq \|f - f_N\|_{L^p}$. Since the triangulation \mathcal{D}_N is a refinement of \mathcal{T}_N^* , we also find that (iii) implies (i) as $\inf_{g \in V_N} \|f - g\|_{L^p} \leq \|f - f_N^*\|_{L^p}$.

We next show that (i) implies (ii). We first note that a consequence of (i) is that

$$\lim_{j \rightarrow +\infty} \sup_{T \in \mathcal{D}_j} e_T(f)_p = 0.$$

It follows that for any $\eta > 0$, there exists $N(\eta)$ such that for $N > N(\eta)$, all triangles $T \in \mathcal{T}_N$ satisfy

$$e_T(f)_p \leq \eta.$$

On the other hand, (i) means that for all $\varepsilon > 0$, there exists $J = J(\varepsilon)$ such that

$$\inf_{g \in V_J} \|f - g\|_{L^p} \leq \varepsilon.$$

For $N > N(\eta)$, we now split \mathcal{T}_N into $\mathcal{T}_N^+ \cup \mathcal{T}_N^-$ where

$$\mathcal{T}_N^+ := \mathcal{T}_N \cap (\cup_{j \geq J} \mathcal{D}_j) \text{ and } \mathcal{T}_N^- := \mathcal{T}_N \cap (\cup_{j < J} \mathcal{D}_j).$$

We then estimate the error of the greedy algorithm by

$$\begin{aligned} \|f - f_N\|_{L^p}^p &= \sum_{T \in \mathcal{T}_N^+} e_T(f)_p^p + \sum_{T \in \mathcal{T}_N^-} e_T(f)_p^p \\ &\leq C^p \sum_{T \in \mathcal{T}_N^+} \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}^p + \eta^p \#(\mathcal{T}_N^-) \\ &\leq C^p \inf_{g \in V_J} \|f - g\|_{L^p}^p + \eta^p \#(\mathcal{T}_N^-) \\ &\leq C^p \varepsilon^p + 2^J N_0 \eta^p, \end{aligned}$$

where C is the stability constant of (2.4). This implies (ii) since for any $\delta > 0$, we can first choose $\varepsilon > 0$ such that $C^p \varepsilon^p < \delta/2$, and then choose $\eta > 0$ such that $2^{J(\varepsilon)} N_0 \eta^p < \delta/2$. When $p = \infty$ the estimate is modified into

$$\|f - f_N\|_{L^\infty} \leq \max\{C\varepsilon, \eta\},$$

which also implies (ii) by a similar reasoning.

We finally prove the equivalence between (i) and (iv) when $p = 2$. Property (i) is equivalent to the L^2 convergence of the orthogonal projection $P_j f$ to f as $j \rightarrow +\infty$, or equivalently of the partial sum

$$\sum_{|\lambda| < j} f_\lambda \psi_\lambda$$

where $|\lambda|$ stands for the scale level of the wavelet ψ_λ . Since $(\psi_\lambda)_{\lambda \in \Lambda}$ is an orthonormal basis of $V(f, \mathcal{R})_2$, the summability and limit of $\sum_{\lambda \in \Lambda} f_\lambda \psi_\lambda$ do not depend on the order of the terms. Therefore (i) is equivalent to the convergence of f_ε to f . \diamond

Remark 3.2 *The equivalence between statements (i) and (iv) can be extended to $1 < p < \infty$ by showing that $(\psi_\lambda)_{\lambda \in \Lambda}$ is an L^p -unconditional system. Recall that this property means that there exists an absolute constant $C > 0$ such that for any finitely supported sequences (c_λ) and (d_λ) such $|c_\lambda| \leq |d_\lambda|$ for all λ , one has*

$$\|\sum c_\lambda \psi_\lambda\|_{L^p} \leq C \|\sum d_\lambda \psi_\lambda\|_{L^p}.$$

A consequence of this property is that if $f \in L^p$ can be expressed as the L^p limit

$$f = \lim_{j \rightarrow +\infty} \sum_{|\lambda| < j} f_\lambda \psi_\lambda,$$

then any rearrangement of the series $\sum f_\lambda \psi_\lambda$ converges towards f in L^p . This easily implies the equivalence between (i) and (iv). The fact that $(\psi_\lambda)_{\lambda \in \Lambda}$ is an unconditional system is well known for Haar systems [19] which correspond to the case $m = 0$, and can be extended to $m > 0$ in a straightforward manner.

3.2 A case of non-convergence

We now show that if we use a greedy refinement rule \mathcal{R} based on a decision function either based on the interpolation or L^2 projection error after bisection, there exists functions $f \in L^p(\Omega)$ such that $f \notin V(f, \mathcal{R})_p$. Without loss of generality, it is enough to construct f on the reference triangle T_{ref} of vertices $\{(0, 0), (1, 0), (1, 1)\}$, since our construction can be adapted to any triangle by an affine change of variables.

Consider first a decision function defined from the interpolation error after bisection, such as (2.5), or more generally

$$d_T(f, e) := \|f - I_{T_e^1} f\|_{L^p(T_e^1)}^p + \|f - I_{T_e^2} f\|_{L^p(T_e^2)}^p.$$

Let f be a continuous function which is not identically 0 on T_{ref} and which vanishes at all points (x, y) such that $x = \frac{k}{2^m}$ for $k = 0, 1, \dots, 2^m$, where m is the degree of polynomial approximation. For such an f , it is easy to see that $I_T f = 0$ and that $I_{T'} f = 0$ for all subtriangles T' obtained by one bisection of T_{ref} . This shows that there is no preferred bisection. Assuming that we bisect from the vertex $(0, 0)$ to the opposite mid-point $(1, \frac{1}{2})$, we find that a similar situation occurs when splitting the two subtriangles. Iterating this observation, we see that an admissible choice of bisections leads after j steps to a triangulation \mathcal{D}_j of T_{ref} consisting of the triangles $T_{j,k}$ with vertices $\{(0, 0), (1, 2^{-j}k), (1, 2^{-j}(k+1))\}$ with $k = 0, \dots, 2^j - 1$. On each of these triangles f is interpolated by the null function and therefore by (2.4) the best L^∞ approximation in V_j does not converge to f as $j \rightarrow +\infty$, i.e. $f \notin V(f, \mathcal{R})_\infty$. It can also easily be checked that $f \notin V(f, \mathcal{R})_p$.

Similar counter-examples can be constructed when the decision function is defined from the L^2 projection error, and has the form

$$d_T(f, e) := \|f - P_{T_e^1} f\|_{L^p(T_e^1)}^p + \|f - P_{T_e^2} f\|_{L^p(T_e^2)}^p.$$

Here, we describe such a construction in the case $m = 1$. We define f on \mathcal{R} as a function of the first variable given by

$$f(x, y) = u(x), \quad \text{if } x \in \left[0, \frac{1}{2}\right], \quad f(x, y) = u\left(x - \frac{1}{2}\right), \quad \text{if } x \in \left(\frac{1}{2}, 1\right],$$

where u is a non-trivial function in $L^2([0, \frac{1}{2}])$ such that

$$\int_0^{\frac{1}{2}} u(x) dx = \int_0^{\frac{1}{2}} xu(x) dx = \int_0^{\frac{1}{2}} x^2 u(x) dx = 0.$$

A possible choice is $u(x) = L_3(4x - 1) = 160x^3 - 120x^2 + 24x - 1$ where L_3 is the Legendre polynomial of degree 3 defined on $[-1, 1]$. With this choice, we have the following result.

Lemma 3.3 *Let T be any triangle such that its vertices have x coordinates either $(0, \frac{1}{2}, 1)$ or $(0, 1, 1)$ or $(\frac{1}{2}, 1, 1)$. Then f is orthogonal to Π_1 in $L^2(T)$.*

Proof: Define

$$T_0 := \left\{ (x, y) \in T, x \in \left[0, \frac{1}{2}\right] \right\} \quad \text{and} \quad T_1 := \left\{ (x, y) \in T, x \in \left(\frac{1}{2}, 1\right] \right\}.$$

Then, with $v(x, y)$ being either the function 1 or x or y , we have

$$\begin{aligned} \int_T f(x, y)v(x, y) dx dy &= \int_{T_0} u(x)v(x, y) dx dy + \int_{T_1} u(x)v(x, y) dx dy \\ &= \int_0^{\frac{1}{2}} u(x)q_0(x) dx + \int_0^{\frac{1}{2}} u(x)q_1(x) dx, \end{aligned}$$

with

$$q_0(x) := \int_{T_{0,x}} v(x, y) dy, \quad q_1(x) := \int_{T_{1,x}} v(x, y) dy,$$

where $T_{i,x} = \{y : (x, y) \in T_i\}$ for $i = 0, 1$. The functions q_0 and q_1 are polynomials of degree at most 2 and we thus obtain from the properties of u that $\int_T f v = 0$. \diamond

The above lemma shows that for any of the three possible choices of bisection of T_{ref} based on the L^2 decision function, the error is left unchanged since the projection of f on all possible sub-triangles is 0. There is therefore no preferred choice, and assuming that we bisect from the vertex $(0, 0)$ to the opposite mid-point $(1, \frac{1}{2})$, then we see that a similar situation occurs when splitting the two subtriangles. The rest of the arguments showing that $f \notin V(f, \mathcal{R})_p$ are the same as in the previous counter-example.

The above two examples of non-convergence reflect the fact that when f has some *oscillations*, the refinement procedure cannot determine the most appropriate bisection. In order to circumvent this difficulty one needs to modify the refinement rule.

3.3 A modified refinement rule

Our modification consists of bisecting from the most recently generated vertex of T , in case the local error is not reduced enough by all three bisections. More precisely, we modify the choice of the bisection of any T as follows:

Let e be the edge which minimizes the decision function $d_T(e, f)$. If

$$(e_{T_e^1}(f)_p^p + e_{T_e^2}(f)_p^p)^{1/p} \leq \theta e_T(f),$$

we bisect T towards the edge e (greedy bisection). Otherwise, we bisect T from its most recently generated vertex (newest vertex bisection). Here θ is a fixed number in $(0, 1)$. In the case $p = \infty$ we use the condition

$$\max\{e_{T_e^1}(f), e_{T_e^2}(f)\} \leq \theta e_T(f).$$

This new refinement rule benefits from the mesh size reduction properties of newest vertex bisection. Indeed, a property illustrated in Figure 1 is that a sequence $\{BNN\}$ of one arbitrary bisection (B) followed by two newest vertex bisections (N) produces triangles with diameters bounded by half the diameter of the initial triangle. A more general property - which proof is elementary yet tedious - is the following: a sequence of the type $\{BNB\cdots BN\}$ of length $k+2$ with a newest vertex bisection at iteration 2 and $k+2$ produces triangles with diameter bounded by $(1-2^{-k})$ times the diameter of the initial triangle, the worst case being illustrated in Figure 2 with $k=3$.

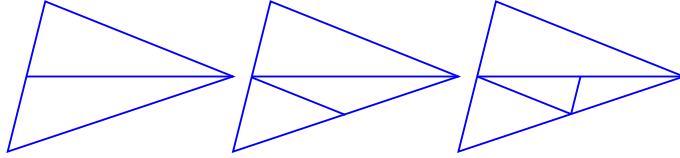


Figure 1: diameter reduction by a $\{BNN\}$ sequence

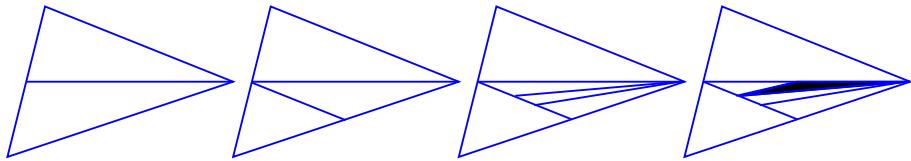


Figure 2: diameter reduction by a $\{BNBBN\}$ sequence
(the dark triangle has diameter at most $7/8$ times the initial diameter)

Our next result shows that the modified algorithm now converges for any $f \in L^p$.

Theorem 3.4 *With \mathcal{R} defined as the modified bisection rule, we have*

$$f \in V(f, \mathcal{R})_p,$$

for any $f \in L^p(\Omega)$ (or $\mathcal{C}(\Omega)$ when $p = \infty$).

Proof: We first give the proof when $p < \infty$. For each triangle $T \in \mathcal{D}_j$ with $j \geq 1$, we introduce the two quantities

$$\alpha(T) := \frac{e_T(f)_p^p}{e_T(f)_p^p + e_{T'}(f)_p^p} \text{ and } \beta(T) := \frac{e_T(f)_p^p + e_{T'}(f)_p^p}{e_{\mathcal{P}(T)}(f)_p^p},$$

where T' is the “brother” of T , i.e. $\mathcal{C}(\mathcal{P}(T)) = \{T, T'\}$. When a greedy bisection occurs in the split of $\mathcal{P}(T)$, we have $\beta(T) \leq \theta^p$. When a newest vertex bisection occurs, we have

$$\beta(T) \leq C^p \frac{\inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}^p + \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T')}^p}{\inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(\mathcal{P}(T))}^p} \leq C^p,$$

where C is the constant of (2.4).

We now consider a given level index $j > 0$ and the triangulation \mathcal{D}_j . For each $T \in \mathcal{D}_j$, we consider the chain of nested triangles $(T_n)_{n=0}^j$ with $T_j = T$ and $T_{n-1} = \mathcal{P}(T_n)$, $n = j, j-1, \dots, 1$. We define

$$\phi\alpha(T) = \prod_{n=1}^j \alpha(T_n) \quad \text{and} \quad \phi\beta(T) = \prod_{n=1}^j \beta(T_n).$$

It is easy to see that

$$\phi\alpha(T)\phi\beta(T) = \frac{e_T(f)_p^p}{e_{T_0}(f)_p^p}$$

so that

$$e_T(f)_p^p \leq C_0 \phi\alpha(T)\phi\beta(T),$$

with $C_0 := \max_{T_0 \in \mathcal{D}_0} e_{T_0}(f)_p^p$. It is also easy to check by induction on j that

$$\sum_{T \in \mathcal{D}_j} \phi\alpha(T) = \sum_{T \in \mathcal{D}_{j-1}} \phi\alpha(T) = \cdots = \sum_{T \in \mathcal{D}_1} \phi\alpha(T) = \#\mathcal{D}_0.$$

We denote by f_j the approximation to f in V_j defined by $f_j = \mathcal{A}_T f$ on all $T \in \mathcal{T}_j$ so that

$$\|f - f_j\|_{L^p}^p = \sum_{T \in \mathcal{D}_j} e_T(f)_p^p.$$

In order to prove that f_j converges to f in L^p , it is sufficient to show that the sequence

$$\varepsilon_j := \max_{T \in \mathcal{D}_j} \min\{\phi\beta(T), \text{diam}(T)\},$$

tends to 0 as j grows. Indeed, if this holds, we split \mathcal{D}_j into two sets \mathcal{D}_j^+ and \mathcal{D}_j^- over which $\phi\beta(T) \leq \varepsilon_j$ and $\text{diam}(T) \leq \varepsilon_j$ respectively. We can then write

$$\begin{aligned} \|f - f_j\|_{L^p}^p &= \sum_{T \in \mathcal{D}_j^+} e_T(f)_p^p + \sum_{T \in \mathcal{D}_j^-} e_T(f)_p^p \\ &\leq C_0 \varepsilon_j \sum_{T \in \mathcal{D}_j^+} \phi\alpha(T) + C^p \sum_{T \in \mathcal{D}_j^-} \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}^p \\ &\leq C_0 \#\mathcal{D}_0 \varepsilon_j + C^p \sum_{T \in \mathcal{D}_j^-} \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^p(T)}^p, \end{aligned}$$

where C is the constant of (2.4). Clearly the first term tends to 0 and so does the second term by standard properties of L^p spaces since the diameter of the triangles in \mathcal{D}_j^- goes to 0. It thus remains to prove that

$$\lim_{j \rightarrow +\infty} \varepsilon_j = 0.$$

Again we consider the chain $(T_n)_{n=0}^j$ which terminates at T , and we associate to it a chain $(q_n)_{n=0}^{j-1}$ where $q_n = 1$ or 2 if bisection of T_n is greedy or newest vertex respectively. If r is the total number of 2 in the chain (q_n) we have

$$\phi\beta(T) \leq C^{pr} \theta^{p(j-r)},$$

with C the constant in (2.4). Let a $k > 0$ be a fixed number, large enough such that

$$C\theta^{k-1} \leq 1.$$

We thus have

$$\phi\beta(T) \leq (C^p \theta^{p(k-1)})^r \theta^{p(j-rk)} \leq \theta^{p(j-rk)}. \quad (3.8)$$

We now denote by l the maximal number of disjoint sub-chains of the type $(\nu_1, 2, \nu_2, \nu_3, \dots, \nu_q, 2)$ with $\nu_j \in \{1, 2\}$ and of length $q+2 \leq 2k+3$ which can be extracted from $(q_n)_{n=0}^{j-1}$. From the remarks on the diameter reduction properties of newest vertex bisection, we see that

$$\text{diam}(T) \leq B(1 - 2^{-2k})^l,$$

with $B := \max_{T_0 \in \mathcal{D}_0} \text{diam}(T_0)$ a fixed constant. On the other hand, it is not difficult to check that

$$r \leq 3l + 3 + \frac{j-r}{2k}. \quad (3.9)$$

Indeed let α_0 be the total number of 1 in the sequence (q_n) which are not preceeded by a 2, and let α_i be the size of the series of 1 following the i -th occurrence of 2 in (q_n) for $i = 1, \dots, r$. Note that some α_i might be 0. Clearly we have

$$j = \alpha_0 + \alpha_1 + \dots + \alpha_r + r.$$

From the above equality, the number of i such that $\alpha_i > 2k$ is less than $\frac{j-r}{2k}$ and therefore there is at least $m \geq r - \frac{j-r}{2k}$ indices $\{i_0, \dots, i_{m-1}\}$ such that $\alpha_i \leq 2k$. Denoting by β_i the position of the i -th occurrence of 2 (so that $\beta_{i+1} = \beta_i + \alpha_i + 1$), we now consider the disjoint sequences of indices

$$S_t = \{\beta_{i_{3t}}, \dots, \beta_{i_{3t+2}}\}, \quad t = 0, 1, \dots$$

There is at least $\frac{m}{3} - 1$ such sequences within $\{1, \dots, j\}$ and by construction each of them contains a sequence of the type $(\nu_1, 2, \nu_2, \nu_3, \dots, \nu_q, 2)$ with $\nu_j \in \{1, 2\}$ and of length $q + 2 \leq 2k + 3$. Therefore the maximal number of disjoint sequences of such type satisfies

$$l \geq \frac{1}{3}(r - \frac{j-r}{2k}) - 1,$$

which is equivalent to (3.9). Therefore according to (3.8)

$$\phi\beta(T) \leq \theta^{p(j-rk)} \leq \theta^{p(\frac{j}{2} - 3(l+1)k + \frac{r}{2})}$$

If $3(l+1)k \leq \frac{j}{4}$, we have

$$\phi\beta(T) \leq \theta^{\frac{pj}{4}}.$$

On the other hand, if $3(l+1)k \geq \frac{j}{4}$, we have

$$\text{diam}(T) \leq B(1 - 2^{-2k})^{\frac{j}{12k} - 1}.$$

We therefore conclude that ε_j goes to 0 as j grows, which proves the result for $p < \infty$.

We briefly sketch the proof for $p = \infty$, which is simpler. We now define $\beta(T)$ as

$$\beta(T) := \frac{e_T(f)_\infty}{e_{\mathcal{P}(T)}(f)_\infty},$$

so that $\beta(T) \leq \theta$ if a greedy bisection occurs in the split of $\mathcal{P}(T)$. With the same definition of $\phi\beta(T)$ we now have

$$e_T(f)_\infty \leq C_0 \phi\beta(T),$$

where $C_0 := \max_{T_0 \in \mathcal{D}_0} e_{T_0}(f)_\infty$. With the same definition of ε_j and splitting of \mathcal{D}_j , we now reach

$$\|f - f_j\|_{L^\infty} \leq \max \left\{ C_0 \varepsilon_j, C \max_{T \in \mathcal{D}_j^-} \inf_{\pi \in \Pi_m} \|f - \pi\|_{L^\infty(T)} \right\},$$

which again tends to 0 if ε_j tends to 0 and f is continuous. The proof that ε_j tends to 0 as j grows is then similar to the case $p < \infty$. \diamond

Remark 3.5 *The choice of the parameter $\theta < 1$ deserves some attention: if it is chosen too small, then most bisections are of type N and we end up with an isotropic triangulation. In the case $m = 1$, a proper choice can be found by observing that when f has C^2 smoothness, it can be locally approximated by a quadratic polynomial $q \in \Pi_2$ with $e_T(f)_p \approx e_T(q)_p$ when the triangle T is small enough. For such quadratic functions $q \in \Pi_2$, one can explicitly study the minimal error reduction which is always ensured by the greedy refinement rule defined a given decision function. In the particular case $p = 2$ and with the choice $\mathcal{A}_T = P_T$ which is considered in the numerical experiments of §4, explicit formulas for the error $\|q - P_T q\|_{L^2(T)}$ can be obtained by formal computing and can be used to prove a guaranteed error reduction by a factor $\theta^* = \frac{3}{5}$. It is therefore natural to choose θ such that $\theta^* < \theta < 1$ (for example $\theta = \frac{2}{3}$) which ensures that bisections of type N only occur in the early steps of the algorithm, when the function still exhibits too many oscillations on some triangles.*

4 Numerical illustrations

The following numerical experiments were conducted with piecewise linear approximation in the L^2 setting: we use the L^2 -based decision function

$$d_T(f, e) := \|f - P_{T_e^1} f\|_{L^2(T_e^1)}^2 + \|f - P_{T_e^2} f\|_{L^2(T_e^2)}^2.$$

and we take for \mathcal{A}_T the $L^2(T)$ -orthogonal projection P_T onto Π_1 . In these experiments, the function f is either a quadratic polynomial or a function with a simple analytic expression which allows us to compute the quantities $e_T(f)_2$ and $d_T(e, f)$ without any quadrature error, or a numerical image in which case the computation of these quantities is discretized on the pixel grid.

4.1 Quadratic functions

Our first goal is to illustrate numerically the optimal adaptation properties of the refinement procedure in terms of triangle shape. For this purpose, we take $f = \mathbf{q}$ a quadratic form i.e. an homogeneous polynomial of degree 2. In this case, all triangles should have the same aspect ratio since the Hessian is constant. In order to measure the quality of the shape of a triangle T in relation to \mathbf{q} , we introduce the following quantity: if (a, b, c) are the edge vectors of T , we define

$$\rho_{\mathbf{q}}(T) := \frac{\max\{|\mathbf{q}(a)|, |\mathbf{q}(b)|, |\mathbf{q}(c)|\}}{|T| \sqrt{|\det(\mathbf{q})|}},$$

where $\det(\mathbf{q})$ is the determinant of the 2×2 symmetric matrix Q associated with \mathbf{q} , i.e. such that

$$\mathbf{q}(u) = \langle Qu, u \rangle$$

for all $u \in \mathbb{R}^2$. Using the reference triangle and an affine change of variables, it is proved in §2 of [12] that

$$e_T(\mathbf{q})_p \sim |T|^{1+\frac{1}{p}} \rho_{\mathbf{q}}(T) \sqrt{|\det(\mathbf{q})|},$$

with equivalence constants independent of \mathbf{q} and T . Therefore, if T is a triangle of given area, its shape should be designed in order to minimize $\rho_{\mathbf{q}}(T)$.

In the case where \mathbf{q} is positive definite or negative definite, $\rho_{\mathbf{q}}(T)$ takes small values when T is isotropic with respect to the metric $|(x, y)|_{\mathbf{q}} := \sqrt{|\mathbf{q}(x, y)|}$, the minimal value $\frac{4}{\sqrt{3}}$ being attained for an equilateral triangle for this metric. Specifically, we choose $\mathbf{q}(x, y) := x^2 + 100y^2$ and display in Figure 3 (left) the triangulation \mathcal{D}_8 obtained after $j = 8$ iterations of the refinement procedure, starting with a triangle which is equilateral for the euclidean metric (and therefore not adapted to \mathbf{q}). Triangles such that $\rho_{\mathbf{q}}(T) \leq 4\sqrt{3}$ (at most 3 times the minimal value) are displayed in white, others in grey. We observe that most triangles produced by the refinement procedure are of the first type and therefore have a good aspect ratio.

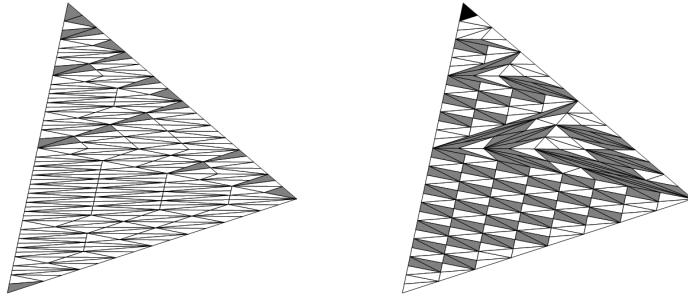


Figure 3: \mathcal{D}_8 for $q(x, y) := x^2 + 100y^2$ (left) and $q(x, y) := x^2 - 10y^2$ (right).

The case of a quadratic function of mixed signature is illustrated in Figure 3 (right) with $\mathbf{q}(x, y) := x^2 - 10y^2$. For such quadratic functions, triangles which are isotropic with respect to the metric $|\cdot|_{\mathbf{q}}$

have a low value of $\rho_{\mathbf{q}}$, where $|\mathbf{q}|$ denotes the positive quadratic form associated to the absolute value $|Q|$ of the symmetric matrix Q associated to \mathbf{q} . Recall for any symmetric matrix Q there exists $\lambda_1, \lambda_2 \in \mathbb{R}$ and a rotation R such that

$$Q = R^T \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} R,$$

and the absolute value $|Q|$ is defined as

$$|Q| = R^T \begin{pmatrix} |\lambda_1| & 0 \\ 0 & |\lambda_2| \end{pmatrix} R.$$

In the present case, $R = I$ and $|\mathbf{q}|(x, y) = x^2 + 10y^2$.

But one can also check that $\rho_{\mathbf{q}}$ is left invariant by any linear transformation with eigenvalues $(t, \frac{1}{t})$ for any $t > 0$ and eigenvectors (u, v) such that $\mathbf{q}(u) = \mathbf{q}(v) = 0$, i.e. belonging to the *null cone* of \mathbf{q} . More precisely, for any such transformation ψ and any triangle T , one has $\rho_{\mathbf{q}}(\psi(T)) = \rho_{\mathbf{q}}(T)$. In our example we have $u = (\sqrt{10}, 1)$ and $v = (\sqrt{10}, -1)$. Therefore long and thin triangles which are aligned with these vectors also have a low value of $\rho_{\mathbf{q}}$. Triangles T such that $\rho_{|\mathbf{q}|}(T) \leq 4\sqrt{3}$ are displayed in white, those such that $\rho_{|\mathbf{q}|}(T) \leq 4\sqrt{3}$ while $\rho_{|\mathbf{q}|}(T) > 4\sqrt{3}$ - i.e. adapted to \mathbf{q} but not to $|\mathbf{q}|$ - are displayed in grey, and the others in dark. We observe that all the triangles produced by the refinement procedure except one are either of the first or second type and therefore have a good aspect ratio.

4.2 Sharp transition

We next study the adaptive triangulations produced by the greedy tree algorithm for a function f displaying a sharp transition along a curved edge. Specifically we take

$$f(x, y) = f_{\delta}(x, y) := g_{\delta}(\sqrt{x^2 + y^2}),$$

where g_{δ} is defined by $g_{\delta}(r) = \frac{5-r^2}{4}$ for $0 \leq r \leq 1$, $g_{\delta}(1 + \delta + r) = -\frac{5-(1-r)^2}{4}$ for $r \geq 0$, and g_{δ} is a polynomial of degree 5 on $[1, 1+\delta]$ which is determined by imposing that g_{δ} is globally C^2 . The parameter δ therefore measures the sharpness of the transition as illustrated in Figure 4. It can be shown that the Hessian of f_{δ} is negative definite for $\sqrt{x^2 + y^2} < 1 + \delta/2$, and of mixed type for $1 + \delta/2 < \sqrt{x^2 + y^2} \leq 2 + \delta$.

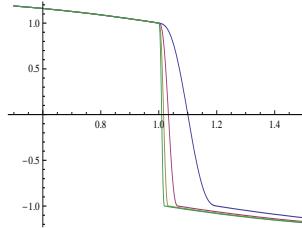


Figure 4: The function g_{δ} , for $\delta = 0.02, 0.03, 0.07, 0.2$.

Figure 5 displays the triangulation \mathcal{T}_{10000} obtained after 10000 steps of the algorithm for $\delta = 0.2$. In particular, triangles T such that $\rho_{\mathbf{q}}(T) \leq 4$ - where \mathbf{q} is the quadratic form associated with $d^2 f$ measured at the barycenter of T - are displayed in white, others in grey. As expected, most triangles are of the first type and therefore well adapted to f . We also display on this figure the adaptive isotropic triangulation produced by the greedy tree algorithm based on newest vertex bisection for the same number of triangles.

Since f is a C^2 function, approximations by uniform, adaptive isotropic and adaptive anisotropic triangulations all yield the convergence rate $\mathcal{O}(N^{-1})$. However the constant

$$C := \limsup_{N \rightarrow +\infty} N \|f - f_N\|_{L^2},$$

strongly differs depending on the algorithm and on the sharpness of the transition, as illustrated in the table below. We denote by C_U , C_I and C_A the empirical constants (estimated by $N \|f - f_N\|_2$ for $N = 8192$) in the uniform, adaptive isotropic and adaptive anisotropic case respectively, and by

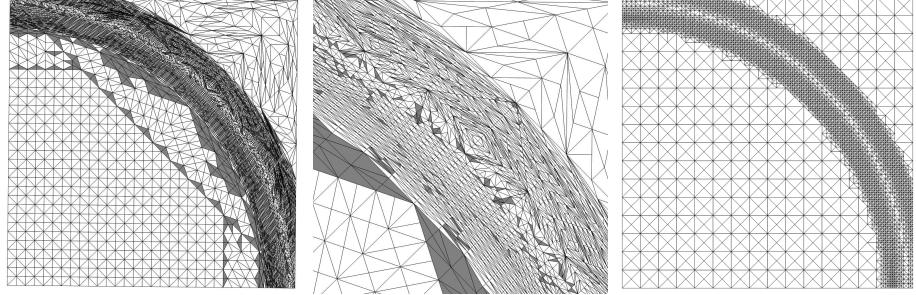


Figure 5: \mathcal{T}_{10000} (left), detail (center), isotropic triangulation (right).

$U(f) := \|d^2f\|_{L^2}$, $I(f) := \|d^2f\|_{L^{2/3}}$ and $A(f) := \|\sqrt{|\det(d^2f)|}\|_{L^{2/3}}$ the theoretical constants suggested by (1.3), (1.2) and (1.1). We observe that C_U and C_I grow in a similar way as $U(f)$ and $I(f)$ as $\delta \rightarrow 0$ (a detailed computation shows that $U(f) \approx 10.37 \delta^{-3/2}$ and $I(f) \approx 14.01 \delta^{-1/2}$). In contrast C_A and $A(f)$ remain uniformly bounded, a fact which reflects the superiority of the anisotropic mesh as the layer becomes thinner.

δ	$U(f)$	$I(f)$	$A(f)$	C_U	C_I	C_A
0.2	103	27	6.75	7.87	1.78	0.74
0.1	602	60	8.50	23.7	2.98	0.92
0.05	1705	82	8.48	65.5	4.13	0.92
0.02	3670	105	8.47	200	6.60	0.92

4.3 Numerical images

We finally apply the greedy tree algorithm to numerical images. In this case the data f has the form of a discrete array of pixels, and the $L^2(T)$ -orthogonal projection is replaced by the $\ell^2(S_T)$ -orthogonal projection, where S_T is the set of pixels with centers contained in T . The approximated 512×512 image is displayed in Figure 6 which also shows its approximation f_{2000} by the greedy tree algorithm based on newest vertex bisection with $N = 2000$ triangles. The systematic use of isotropic triangles results in strong ringing artifacts near the edges.

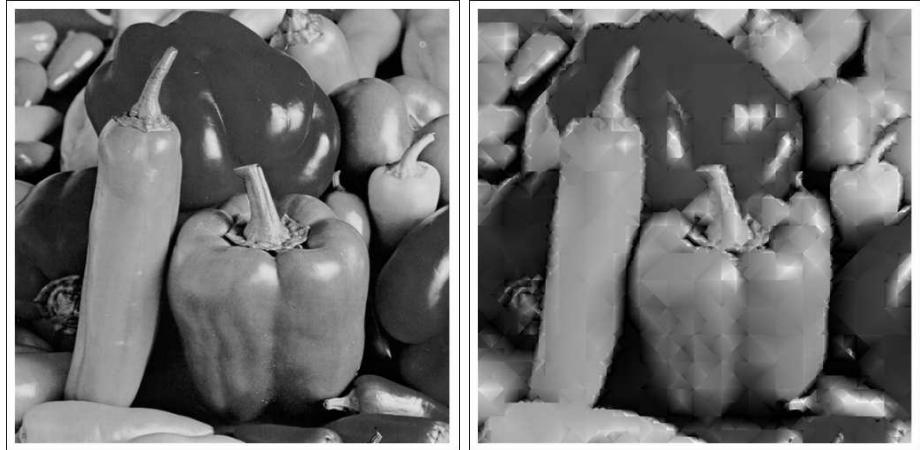


Figure 6: The image "peppers" (left), f_{2000} with newest vertex (right).

We display in Figure 7 the result of the same algorithm now based on our greedy bisection procedure with the same number of triangles. As expected, the edges are better approximated due to the presence of well oriented anisotropic triangles. Yet artifacts persist on certain edges due to oscillatory features in the image which tend to mislead the algorithm in its search for triangles with good aspect ratio, as explained in §3.2. These artifacts tend to disappear if we use the modified refinement rule proposed in §3.3 as also

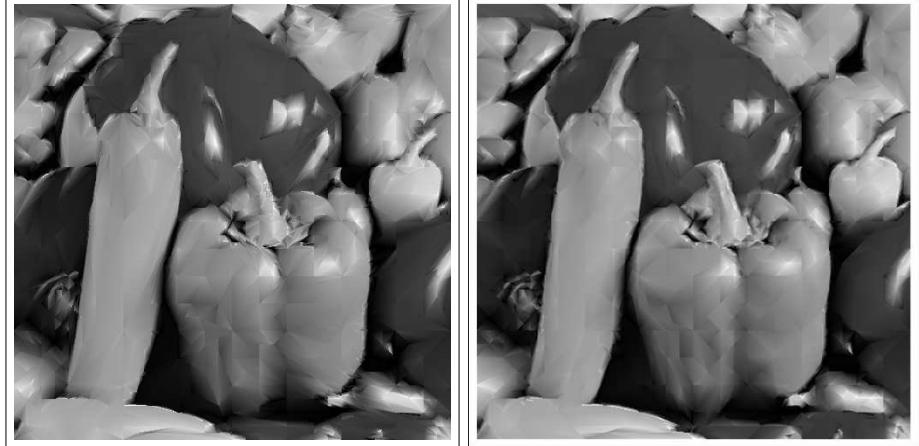


Figure 7: f_{2000} with greedy bisection (left), modified procedure (right).

illustrated on Figure 7. This modification is thus useful in the practical application of the algorithm, in addition of being necessary for proving convergence convergence of the approximations towards any L^p function. Note that encoding a triangulation resulting from N iterations of the anisotropic refinement algorithm is more costly than for the newest vertex rule: the algorithm encounters at most $2N$ triangles and for each of them, one needs to encode one out of four options (bisect towards edge a or b or c or not bisect), therefore resulting into $4N$ bits, while only two options need to be encoded when using the newest vertex rule (bisect or not), therefore resulting into $2N$ bits. In the perspective of applications to image compression, another issue is the quantization and encoding of the piecewise affine function as well as the treatment of the triangular visual artifacts that are inherent to the use of discontinuous piecewise polynomials on triangulated domains. These issues will be discussed in a further work specifically dealing with image applications.

5 Conclusions and perspectives

In this paper, we have studied a simple greedy refinement procedure which generates triangles that tend to have an *optimal aspect ratio*. This fact is rigorously proved in [12], together with the optimal convergence estimate (1.1) for the adaptive triangulations constructed by the greedy tree algorithm in the case where the approximated function f is C^2 and convex. Our numerical results illustrate these properties.

In the present paper we also show that for a general $f \in L^p$ the refinement procedure can be misled by oscillations in f , and that this drawback may be circumvented by a simple modification of the refinement procedure. This modification appears to be useful in image processing applications, as shown by our numerical results.

Let us finally mention several perspectives that are raised from our work, and that are the object of current investigation:

1. Conforming triangulations: our algorithm inherently generates hanging nodes, which might not be desirable in certain applications, such as numerical discretization of PDE's where anisotropic elements are sometimes used [3]. When using the greedy tree algorithm, an obvious way of avoiding this phenomenon is to bisect the chosen triangle together with an adjacent triangle in order to preserve conformity. However, it is no more clear that this strategy generates optimal triangulations. In fact, we observed that many inappropriately oriented triangles can be generated by this approach. An alternative strategy is to apply the non-conforming greedy tree algorithm until a prescribed accuracy is met, followed by an additional refinement procedure in order to remove hanging nodes.
2. Discretization and encoding: our work is in part motivated by applications to image and terrain data processing and compression. In such applications the data to be approximated is usually given in discrete form (pixels or point clouds) and the algorithm can be adapted to such data, as shown

in our numerical image examples. Key issues which need to be dealt with are then (i) the efficient encoding of the approximations and of the triangulations using the tree structure in a similar spirit as in [11] and (ii) the removal of the triangular visual artifacts due to discontinuous piecewise polynomial approximation by an appropriate post-processing step.

3. Adaptation to curved edges: one of the motivation for the use of anisotropic triangulations is the approximation of functions with jump discontinuities along an edge. For simple functions, such as characteristic functions of domains with smooth boundaries, the L^p -error rate with an optimally adapted triangulation of N elements is known to be $\mathcal{O}(N^{-\frac{2}{p}})$. This rate reflects an $\mathcal{O}(1)$ error concentrated on a strip of area $\mathcal{O}(N^{-2})$ separating the curved edge from a polygonal line. Our first investigations in this direction indicate that the greedy tree algorithm based on our refinement procedure cannot achieve this rate, due to the fact that bisection does not offer enough geometrical adaptation. This is in contrast with other splitting procedures, such as in [13] in which the direction of the new cutting edge is optimized within an infinite range of possible choices, or [15] where the number of choices grows together with the resolution level. An interesting question, addressed and partially answered in §9.1 of [24], is thus to understand if the optimal rate for edges can be achieved by a splitting procedure with a small and fixed number of choices similar to our refinement procedure, which would be beneficial from both a computational and encoding viewpoint.

References

- [1] F. Arandiga, A. Cohen, R. Donat, N. Dyn and B. Matei, *Approximation of piecewise smooth images by edge-adapted techniques*, ACHA 24, 225–250, 2008.
- [2] B. Alpert, *A class of bases in L^2 for the sparse representation of integral operators*, SIAM J. Math. Anal. 24, 246–262, 1993.
- [3] T. Apel, *Anisotropic finite elements: Local estimates and applications*, Series “Advances in Numerical Mathematics”, Teubner, Stuttgart, 1999.
- [4] V. Babenko, Y. Babenko, A. Ligun and A. Shumeiko, *On Asymptotical Behavior of the Optimal Linear Spline Interpolation Error of C^2 Functions*, East J. Approx. 12(1), 71–101, 2006.
- [5] P. Binev, W. Dahmen and R. DeVore, *Adaptive Finite Element Methods with Convergence Rates*, Numerische Mathematik 97, 219–268, 2004.
- [6] R. Baraniuk, H. Choi, J. Romberg and M. Wakin, *Wavelet-domain approximation and compression of piecewise smooth images*, IEEE Transactions on Image Processing, 15(5), 1071–1087, 2006.
- [7] H. Borouchaki, P.J. Frey, P.L. George, P. Laug and E. Saltel, *Mesh generation and mesh adaptivity: theory, techniques*, in Encyclopedia of computational mechanics, E. Stein, R. de Borst and T.J.R. Hughes ed., John Wiley & Sons Ltd., 2004.
- [8] L. Breiman, J.H. Friedman, R.A. Olshen and C.J. Stone, *Classification and regression trees*, Wadsworth international, Belmont, CA, 1984.
- [9] E. Candes and D. L. Donoho, *Curvelets and curvilinear integrals*, J. Approx. Theory. 113, 59–90, 2000.
- [10] L. Chen, P. Sun and J. Xu, *Optimal anisotropic meshes for minimizing interpolation error in L^p -norm*, Math. of Comp. 76, 179–204, 2007.
- [11] A. Cohen, W. Dahmen, I. Daubechies and R. DeVore, *Tree-structured approximation and optimal encoding*, App. Comp. Harm. Anal. 11, 192–226, 2001.
- [12] A. Cohen and J.M. Mirebeau, *Greedy bisection generates optimally adapted triangulations*, Accepted in Mathematics of Computation, 2010.

- [13] S. Dekel and D. Leviathan, *Adaptive multivariate approximation using binary space partitions and geometric wavelets*, SIAM Journal on Numerical Analysis 43, 707–732, 2005.
- [14] L. Demaret, N. Dyn, M. Floater and A. Iske, *Adaptive thinning for terrain modelling and image compression*, in Advances in Multiresolution for Geometric Modelling, N.A. Dodgson, M.S. Floater, and M.A. Sabin (eds.), Springer-Verlag, Heidelberg, 321-340, 2005.
- [15] D. Donoho, *Wedgelets: nearly minimax estimation of edges*, Ann. Statist. 27(3), 859–897, 1999.
- [16] D. Donoho, *CART and best basis: a connexion*, Ann. Statist. 25(5), 1870–1911, 1997.
- [17] W. Dörfler, *A convergent adaptive algorithm for Poisson’s equation*, SIAM J. Numer. Anal. 33, 1106–1124, 1996.
- [18] B. Karaivanov and P. Petrushev, *Nonlinear piecewise polynomial approximation beyond Besov spaces*, Appl. Comput. Harmon. Anal. 15(3), 177-223, 2003.
- [19] B. S. Kashin and A. A. Saakian, *Orthogonal series*, Amer. Math. Soc., Providence, 1989.
- [20] E. Le Pennec and S. Mallat, *Bandelet image approximation and compression*, SIAM Journal of Multiscale Modeling. and Simulation, 4(3), 992–1039, 2005.
- [21] P. Morin, R. Nochetto and K. Siebert, *Convergence of adaptive finite element methods*, SIAM Review 44, 631–658, 2002.
- [22] R. Stevenson, *An optimal adaptive finite element method*, SIAM J. Numer. Anal., 42(5), 2188–2217, 2005.
- [23] R. Verfurth, *A Review of A Posteriori Error Estimation and Adaptive Mesh-Refinement Techniques*, Wiley-Teubner, 1996.
- [24] J.-M. Mirebeau, *Adaptive and anisotropic finite element approximation : Theory and algorithms*, PhD Thesis, tel.archives-ouvertes.fr/tel-00544243/en/

Albert Cohen

UPMC Univ Paris 06, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 cohen@ann.jussieu.fr

Nira Dyn

School of Mathematics, Tel Aviv University, Ramat Aviv, Israel
 niradyn@math.tau.ac.il

Frédéric Hecht

UPMC Univ Paris 06, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 hecht@ann.jussieu.fr

Jean-Marie Mirebeau

UPMC Univ Paris 06, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France
 mirebeau@ann.jussieu.fr